CONTENT IDENTIFICATION AND MANAGEMENT AGREEMENT

RIGHTS OWNER FULL LEGAL NAME: **Viacom Inc. (and its wholly-owned affiliate entities)** (hereinafter the "Rights Owner")	**TYPE OF ENTITY:** [X] Corporation [ ] Limited Liability Company [ ] Sole Proprietorship [ ] Other [specify _____]
COUNTRY (AND STATE IF IN THE UNITED STATES) OF INCORPORATION OR RESIDENCE: New York	**TAX IDENTIFICATION NUMBER:** N/A

	BUSINESS CONTACT	TECHNICAL CONTACT	ACCOUNTING CONTACT
Name:	Stanley Pierre-Louis	Alan Bell	
Title:	VP, Associate General Counsel	Executive VP, CTO	
Address: **City, State:** **Postal Code:** **Country:**	Viacom, Inc. 1515 Broadway New York, NY 10036	Paramount Pictures Corporation 5555 Melrose Avenue Hollywood, CA 90038	
Phone:	(212) 846-4811	(323) 956-8990	
Fax:	(201) 553-7714	(818) 571-1335	
Email:	stanley.pierre-louis@viacom.com	Alan_Bell@Paramount.com	

CONTENT LICENSE OPTION:
[_] Rights Owner agrees to license and monetize content pursuant to the Content License Agreement.
[_] Rights Owner agrees to license and monetize content pursuant to a separate license agreement with Google dated [_____] titled [_____].
[X] Rights Owner does not agree to license and monetize content, and elects only to block or track content.

Rights Owner and Google hereby agree to this Content Identification and Management Agreement ("Agreement").

Effective Date: February 1, 2008

Google Inc.

BY: David H. Eun
DAVID EUN
Vice President, Content Partnerships
Google, Inc.
NAME: _____
TITLE: _____
1600 Amphitheatre Parkway
Mountain View, CA 94043

2008.02.20
10:45:55
-08'00'

Viacom Inc.
(and its wholly-owned affiliate entities)

BY: _____
NAME: Michael D. Fricklas
TITLE: Executive VP, General Counsel
1515 Broadway
New York, NY 10036

Google's content identification and management system ("System") and content preparation software ("Software") are designed to help Rights Owner identify its Works on YouTube and set Usage Policies for those Works. The following terms govern Rights Owner's use of the System and Software.

1. Definitions.

"Block" means the Usage Policy available in the System for Rights Owner to specify that a user video be blocked from playback on YouTube in the territories selected by Rights Owner.
"ID File" means the unique binary data that describes a Work and is used for the automatic identification of that Work or a portion thereof. ID Files may be provided by Rights Owner to Google or created by Google using the Reference Files.
"Monetize" means the Usage Policy available in the System for Rights Owner to license to Google in the territories selected by Rights Owner a user video matching an ID File or claimed by Rights Owner using the search functionality that may be offered by the System.
"Reference Files" means the Works provided by Rights Owner to Google pursuant to this Agreement.
"Software" has the meaning given in the preamble.
"System" has the meaning given in the preamble.
"Track" means the Usage Policy available in the System for Rights Owner where it does not specify that the user video be blocked from playback on YouTube, but also does not grant any licenses thereto.
"Usage Policy" means Monetize, Track, or Block, or such other policies as may be made available by Google from time to time.
"Work" means audio and audiovisual works owned or controlled by Rights Owner.
"YouTube" means YouTube.com and subdomains.

2. Reference Files and ID Files. (a) Rights Owner will deliver to Google the Reference Files or ID Files created using the Software. If Rights Owner provides Reference Files, Google will create corresponding ID Files. Rights Owner shall retain all rights, including without limitation copyright rights, in Reference Files. Rights Owner will provide metadata associated with each Reference File or ID File (such as title, description, the Usage Policy, and the territorial scope of each Usage Policy) via an XML feed or otherwise pursuant to Google's reasonable specifications. Rights Owner will make commercially reasonable efforts to ensure that the metadata delivered to Google is accurate and current. Google will provide appropriate format, resolution, and bit-rate specifications for the delivery of Reference Files, ID Files, and metadata. (b) Rights Owner may inactivate from use in the System any of its Reference Files and ID Files at any time and thereby terminate Google's license to use the Reference Files and ID Files. Google will store the Reference Files and ID Files on secure servers and will protect Reference Files and ID Files from unauthorized access as specified in Exhibit A. Google will develop the capability to delete or destroy, at Rights Owner's Request, any or all of Rights Owner's Reference Files and ID Files; provided, however, that nothing herein alters either party's document retention or discovery obligations in connection with any pending or future litigation between the parties, and Google's retention of ID Files or Reference Files in compliance with any such obligations will not be deemed a breach of this Agreement. Google will use commercially reasonable efforts to implement such capability no later than July 31, 2008.

3. User Video Matches. The System will compare all videos uploaded to YouTube, including all videos designated "private" and those available through versions of YouTube localized for particular countries, against the ID Files to identify matches and apply the Usage Policies assigned by Rights Owner to any matches. Google will use commercially reasonable efforts to improve the System with the goal of minimizing the time between video upload and application of the Usage Policies set by Rights Owner. The System may also provide Rights Owner the capability to perform text searches for user videos that may contain the Works and assign Usage Policies for such materials. Rights Owner may change any Usage Policy at any time. If a particular ID File has not yielded any matches within a one-year period of time, Google may by written notice request from Rights Owner permission to remove such ID file from the System, which Rights Owner may authorize in its sole discretion. Google may replace old ID Files with new ID Files of a particular work at any time in connection with System upgrades and technical

modifications. Rights Owner shall not knowingly make false claims on user videos. Knowingly false claims may lead to termination of this Agreement by Google.

4. Disputes. Google may establish reasonable procedures to resolve claims that appear to be in good faith by a user that a Work has been blocked due to error, mistake, or otherwise. Rights Owner will cooperate with Google to resolve such disputes. If, during the course of evaluating such claims, Rights Owner reviews content designated as private by the user, Rights Owner will not disclose the content to any third party except as necessary to complete its evaluation process or in contemplation of, or participation in, a judicial proceeding. Notwithstanding the foregoing, nothing herein shall limit Rights Owner's rights and remedies under applicable law against a user with respect to any video under review.

5. Licenses and Ownership. (a) Google grants to Rights Owner a non-exclusive, non-transferable, royalty-free, limited license to use the System and Software solely for the purpose of creating ID Files and identifying and managing its Works on YouTube. By providing Reference Files and/or ID Files, Rights Owner grants Google a non-exclusive, non-transferable, royalty-free, limited license to store, copy (including the right to make temporary cache and storage copies), modify or reformat, excerpt, analyze, use to create algorithms and binary representations, and otherwise use those files solely in connection with the System and subject to the terms of this Agreement. **(b)** Rights Owner shall not sell, lease, lend, convey, modify, adapt, translate, prepare derivative works from, decompile, reverse engineer, disassemble or attempt to derive source code from the System or Software. All rights or licenses not explicitly granted by the parties herein are specifically reserved. Except for the licenses specifically granted above, all of Rights Owner's intellectual property rights in the Reference Files and ID Files (whether provided by Rights Owner to Google or created by Google) remain with Rights Owner, and all of Google's intellectual property rights in YouTube, the Software, the System and related information and files remain with Google. For the avoidance of doubt, Rights Owner does not grant Google the right to modify, adapt, prepare derivative works, store or reproduce Reference Files and ID Files except as necessary to comply with the terms of this Agreement, nor does Rights Owner grant Google the right to publicly perform, publicly display, or distribute Reference Files and ID Files. Upon any termination of this Agreement, both parties will delete all ID Files from their respective storage systems.

6. Confidentiality. Neither party will disclose the terms of this Agreement to any third party (except to outside counsel or retained experts), or issue any public announcement regarding the terms of this Agreement, without the other party's prior written agreement. The parties shall not disclose to any third parties nonpublic information disclosed by one party to the other under this Agreement, and shall protect such information applying the same degree of care used by the parties to protect their own confidential information. If this Agreement or any confidential information of either party is required to be produced by law, the noticed party will promptly notify the other party and, to the extent practicable, cooperate to obtain an appropriate protective order prior to disclosing any confidential information. Except with respect to the terms and existence of this Agreement, this Agreement imposes no obligation upon Google or Rights Owner with respect to the other party's confidential information that: (i) a party knew before receiving it from the other party pursuant to this Agreement or a party knew before participating in the System; (ii) becomes publicly available through no fault of the other party; (iii) is rightfully received by the other party from a third party without a duty of confidentiality; or (iv) is independently developed without reference to Google's confidential information.

7. Representations and Warranties, Indemnities. Each party represents and warrants that it has authority to grant the licenses set forth in Section 5. Rights Owner represents and warrants that it believes in good faith, after reasonable investigation, that it has all rights required to set the Usage Policies that it has set with respect to its Works. Each party shall indemnify, defend and hold harmless the other party, and their respective directors, officers, employees, and agents from any third party claims arising out of a breach of that party's representations and warranties.

8. DISCLAIMERS, LIMITATIONS OF LIABILITY. EXCEPT FOR THE EXPRESS WARRANTIES MADE BY THE PARTIES IN SECTION 7, THE PARTIES DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. EXCEPT FOR THE INDEMNIFICATION OBLIGATIONS IN SECTION 7, NEITHER PARTY WILL BE LIABLE TO THE OTHER FOR INDIRECT, CONSEQUENTIAL, SPECIAL, PUNITIVE OR EXEMPLARY DAMAGES OR PENALTIES ARISING FROM ANY ACTION TAKEN PURSUANT TO THIS AGREEMENT. PRIOR TO RIGHTS OWNER PROVIDING REFERENCE FILES TO GOOGLE FOR THE PREPARATION OF ID FILES, THE PARTIES AGREE TO ENTER INTO GOOD FAITH NEGOTIATIONS LIMITING GOOGLE'S AGGREGATE LIABILITY FOR ANY CAUSE OF ACTION ARISING FROM OR RELATED TO BREACHES OF THE SECURITY PROVISIONS IN EXHIBIT A RESULTING IN A REFERENCE FILE BEING WRONGFULLY COPIED OR ACQUIRED BY ANY THIRD PARTY. FOR THE AVOIDANCE OF DOUBT, NOTHING HEREIN SHALL BE DEEMED A RELEASE OR WAIVER BY RIGHTS OWNER WITH RESPECT TO CLAIMS FOR DAMAGES ARISING FROM THE PRESENCE OF A WORK ON YOUTUBE THAT HAS NOT BEEN LICENSED TO GOOGLE BY RIGHTS OWNER; PROVIDED, HOWEVER, GOOGLE SHALL NOT BE LIABLE TO RIGHTS OWNER FOR ANY AMOUNT UNDER ANY THEORY OF LIABILITY WITH RESPECT TO THOSE ID FILES FOR WHICH RIGHTS OWNER AFFIRMATIVELY ELECTS THE "TRACK" USAGE POLICY.

9. NO EFFECT ON PENDING, FUTURE, OR RELATED LITIGATION. Notwithstanding the foregoing, nothing in this Agreement shall limit or expand in any way whatsoever Google's and/or Rights Owner's pursuit or introduction of evidence in any litigation or contemplated litigation between them, including but not limited to *Viacom International, Inc. et al v. YouTube, Inc. et al.*, Case No. 1:07-cv-02103-LLS, filed on March 13, 2007, and currently pending in the United States District Court for the Southern District of New York. Furthermore, nothing in this Agreement shall be cited as a defense against or agreement to the production of any relevant material in discovery in any lawsuit, subject to any protective order entered in such lawsuit, and nothing in this Agreement shall operate in any respect as a release or waiver of any of the claims in any lawsuit except as expressly provided in Section 8.

10. Termination. (a) Either party may end this Agreement on 30 days written notice. All licenses granted in this Agreement will expire upon termination. **(b)** Sections 1, 5(b), 6-8, 9(b), and 10 survive termination.

11. Miscellaneous. The parties are independent contractors, and nothing in this Agreement creates an agency, partnership, or joint venture. Neither party may assign rights or obligations under this Agreement to any third party without the prior written consent of the other. This Agreement sets forth the entire agreement between the parties and supersedes any prior or contemporaneous agreements regarding its subject matter. This Agreement may be amended only in a writing signed by both parties. If this Agreement conflicts with any other agreement applying to Google's use of Works on YouTube, these terms control. Each party will send any notices hereunder in writing and to the attention of the Legal Department at the address listed on the first page of this Agreement. If any provision of this Agreement conflicts with applicable laws or is adjudicated to be illegal, that provision will be deemed eliminated from the Agreement and the Agreement will remain in effect so long as the essential purpose can still be achieved. This Agreement is governed by the laws of the State of California (excluding its choice of law rules) and applicable federal laws. Except with respect to claims or actions involving users pursuant to Section 4, any litigation to enforce the terms of this Agreement will be brought in any state or federal court of competent jurisdiction in Santa Clara County, California; each party consents to venue and exclusive personal jurisdiction of such courts. This Agreement may be executed in one or more counterparts, each of which will be deemed an original and all of which, when taken together, will constitute a single instrument.

EXHIBIT A
SECURITY DOCUMENT

Security Overview for Video ID at Google

Introduction

Securing networked servers against would-be hackers is key to ensuring the success of any system. When it comes to partner collaboration, the importance is paramount. Google invests billions of dollars in technology, people, and process to ensure data at Google is safe, secure, and private. Google's dedicated team of security professionals is responsible for designing in security from the onset, reviewing all design, code, and finished product to ensure it meets strict Google security and data privacy standards. The same infrastructure used to host various applications at Google and to secure hundreds of thousands of user's data is also used to manage millions of consumers' data and billions of dollars in advertising transactions. Customer information and files are safe and secure.

This document describes the security measures and controls that Google has put in place to ensure the security of customer data. The key aspects covered include:

- Physical security and internal information security at Google data centers
- Change management processes, data backup/destruction, privacy policy
- Infrastructure for Video ID data

This document describes a snapshot of the current procedures for security. Google reserves the right to adjust these measures as our systems change and attackers adapt.

Security Team

Google employs a large team of information security experts to design and maintain our defense systems, and to make security a core part of the development philosophy and culture. Because we must protect the data of hundreds of millions of end users, we take extra care to make sure that all applications and services that we launch are secure.

Google's security team consists of some of the most accomplished security veterans in the IT industry. Many have experience running security operations at Fortune 500 companies, including some of the most well known financial service institutions. Examples of the backgrounds of individuals on the security team include:

- Chief Information Security Officer at Charles Schwab
- Director of Secure Networking Research at Bell Labs
- Technical Director for Information Security at Charles Schwab
- Senior Network Forensics Specialist from the National Nuclear Security Administration

The security team is involved in all aspects of the security process at Google, including the construction of a custom security infrastructure tuned to Google's unique platforms. They are responsible for the perimeter defense systems described below, as well as the security review process for applications described later in the document.

Data Center Environment and Physical Security

Google Data Center Infrastructure

Google maintains a vast number of geographically distributed data centers located primarily in the USA and the European Union. Data centers are unmarked and in undisclosed locations to maximize security.

Physical Security Staffing

At the Google data centers, there is a Security Operations Center, which is manned 24 hours a day, 7 days a week by a physical security services organization. The security organization deploys three shifts of 8 hours to provide continuous coverage. The security operations centers contain the monitors for the Closed Circuit TV (CCTV) cameras and all alarm systems. Internal and external patrols of the data center are performed each shift. The data centers are housed in facilities that require electronic key access, with alarms that are linked to the guard station manned 24 hours a day, 7 days a week.

Physical Security Access Procedures

Formal access procedures exist for allowing physical access to the data centers. All entrants to the data center must identify themselves as well as show proof of identity to security operations. Valid proof of identity is a photo ID issued by Google and a governmental entity. Only authorized Google employees and contractors are allowed entry to the data centers. Data center managers must approve any visitors in advance for the specific data center and internal areas they wish to visit.

Only authorized Google employees and contractors who permanently work at the data centers are permitted to request card access to these facilities. Data center card access requests must be made through e-mail, and requires the approval of the requestor's manager and the Data Center Director. All other Google employees and authorized contractors requiring temporary data center access must sign in at the guard station, present an Google badge (Google employees or contractors) or ID issued by their employer (authorized contractors) and reference an approved data center access record identifying the individual as approved.

Physical Security Devices

The data centers employ electronic card key and biometric access control system that are linked to a system alarm. The access control system monitors and records each individual's access to perimeter doors, shipping/receiving, the raised floor, and other critical areas. Unauthorized activity and failed access attempts are logged by the access control system, investigated as appropriate, and reported to the security manager. The security manager reviews and approves these reports. Authorized access throughout the business operations and data centers is restricted based on zones and the individual's job responsibilities.

All entrants to the data centers must pass through a mantrap. The mantrap is designed to physically limit access to one person at a time (floor sensors and automatic 180 degree turnstile) and prohibits the "handing off" of a badge back to a secondary person.

The fire doors at the data centers are alarmed and can only be opened from the inside. The fire doors are fitted with push bars to open. There is a specified delay on the push bar unless a fire alarm has been activated. If a person tries to exit the building through a fire door without a fire alarm having been triggered, an alarm would register in the security operations center.

CCTV cameras are in operation both inside and outside the data centers. The positioning of the cameras has been designed to cover strategic areas including, among others, the perimeter, doors to the data center building, shipping/receiving and the raised floor.

Security operations personnel manage the CCTV monitoring, recording and control equipment. The CCTV equipment is connected by secure cables throughout the data centers. Cameras record on site via digital video recorders 24 hours a day, 7 days a week. The surveillance records are retained for 60-90 days based on activity.

Environmental Safeguards

Redundancy

The data centers are designed with resiliency and redundancy. The redundancy is intended to minimize the impact of common equipment failures and environmental risks. Infrastructure systems have been designed to eliminate single points of failure. Dual circuits, switches, networks or other necessary devices are utilized to provide this redundancy. Critical facilities infrastructure at the data centers have been designed to be robust, fault tolerant and concurrently maintainable. Preventative and corrective maintenance is performed without interruption of services.

All environmental equipment and facilities have documented preventative maintenance procedures that detail the procedure and frequency of performance in accordance with the manufacturer's or internal specifications. Preventative and corrective maintenance of the Google data center equipment is scheduled through the standard change process. Preventative maintenance is performed on all infrastructure systems according to documented procedures.

Power

The data center electrical power systems are designed to be fully redundant and maintainable without impact to continuous operations, 24 hours a day, and 7 days a week. A primary as well as an alternate power source, each with equal capacity, is provided for every critical infrastructure component in the data center. This redundancy begins with dual utility power feeds, primary and alternate, to parallel utility switchboards sized so that any one can provide power to the entire facility. The output power is then routed to Automated Transfer Switches (ATS), which supply all building loads including uninterruptible power supplies (UPS), building and mechanical services, and heating, ventilation and air conditioning systems.

Battery backup power is provided by UPS batteries, which supply consistently reliable power protection during utility brownouts, blackouts, over voltage, under voltage, and out-of-tolerance frequency conditions. During normal operations, the utility power charges the battery modules as well as supplies power to the data center raised floor. If utility power is interrupted, the UPS batteries provide back-up until the diesel generator systems take over. In the event of unavailability of both electrical utility and diesel generators, the UPS batteries can provide emergency electrical power to run the data center at full capacity for 10 minutes.

If utility power is interrupted or is out of specification, the power supply will automatically switch to battery mode to continue to supply power to the data center without interruption. When utility power returns, the switch will remain in bypass so that the data center operations team can ascertain the issue has been corrected and can bring the systems back to normal mode in a controlled manner.

Solid State Transition Transfer Switches (SSTTS) are also in place. Should UPS power fail, the SSTTS can be used to transparently transfer all loads from the external dual utility power feeds to the diesel generators.

Diesel engine generators are in place to provide power to critical equipment and customer loads. The generators are capable of providing enough emergency electrical power to run the data center at full capacity typically for a period of days. These generators automatically startup and provide power within seconds in the event of a power outage.

The automatic startup and power distribution is controlled by a programmable logic controller, which has a redundant backup.

Google has short notice diesel refueling contracts in place.

Climate and Temperature

Air-cooling is required to maintain a constant operating temperature for servers and other computing hardware, which prevents over heating and reduces the possibility of service outage. Computer room air conditioning units are powered by both normal and emergency electrical systems. Security operations teams monitor these units and perform periodic inspections and preventative maintenance.

Fire Detection and Suppression

At the data center, automatic fire detection and suppression equipment has been installed to prevent damage to computing hardware. The fire detection systems utilize heat, smoke, and/or water detection sensors that are located in the data center ceiling as well as underneath the raised floor.

In the event fire or smoke is detected, the detection system will sound audible and/or visible alarms in the zone affected, at the security operations console and at the remote monitoring desk of the local fire department. .

In addition, there are fire extinguishers located throughout the data centers.

Logical Software Infrastructure Security Measures

Google Server Environment

Google's servers are designed in-house from the ground up, and Google maintains control over the entire hardware and software stack. The operating system is based on Linux, and has been customized to solely run Google services. Since these are not meant to be general purpose systems like a typical OS, the core services and binaries of the OS have been stripped down, hardened, and heavily modified to leave only those necessary to run Google's applications.

As a result of this degree of control and homogeneity over the entire stack, Google designs its security infrastructure in a very different way from traditional systems. Rather than having to guard against a wide array of unknown inputs into many third party applications, Google can anticipate exactly what types of queries can come into the system and only accept this whitelisted set of queries. This philosophy is utilized throughout the security framework to only accept what is expected, and this provides a highly secure application environment.

Firewalls and Intrusion Detection

Google employs multiple layers of firewalls and intrusion detection to ensure that that our external attack surface is protected.

Intrusion detection is intended to provide insight into ongoing attack activities and provide adequate information to respond to incidents. Many companies make extensive use of third-party technologies (e.g., Network Intrusion Detection Systems - NIDS, Host-based Intrusion Detection Systems - HIDS) to look for known attacks against commonly-installed software, and Security Operations Centers (SOCs) to respond when they arise.

We take a different approach by:

- Tightly controlling the size and make-up of our attack surface through preventative measures
- Employing intelligent detect controls at data entry points
- Employing technologies that automatically remedy dangerous situations.

Most of our Internet-exposed attack surface is comprised of Google-created software and the internal environment is large and complex. Traditional IDS products are not economical, efficient or useful in these situations and we have needed to rely on smarter methods of detecting exploitation.

When we approach intrusion detection concepts, we break down our attack surface according to anticipated input vectors (i.e., how hackers will attempt to break in). All of Google's hosting infrastructure is custom-built so we have the ability to tightly define our perimeter and the entrance points into our network.

While we cannot talk in detail about the exact defenses without potentially compromising Google's defense system, some of the major areas of coverage that achieve the goals above are as follows:

- As mentioned previously, the OS on every system is stripped down, modified, and hardened to avoid third party vulnerabilities on running systems
- All IP traffic is routed through custom front end servers that detect and stop malicious requests
- Traffic is examined for exploitation of programming errors via methods such as cross-site scripting, and high priority alerts are generated when such an exploit is found
- To prevent buffer overflow attacks, all open source software that is Internet facing or that processes external data goes through several levels of code review, audit, and modification before allowed into production. All changes are contributed back to the open source community.
- Systems are checked continually for binary modifications, and any unrecognized modifications are purged
- Router ACLs are used to provide perimeter defense, and an internally routable IP space is used to make sure external connections are never made to internal systems
- Layer 3 filtering is used to mitigate packet-level attacks

Multi-tenant Distributed Data Environment

Google applications run in a multi-tenant distributed environment. Rather than segregating customer data to one machine or set of machines, data from all customers is

distributed amongst a shared infrastructure of tens of thousands of homogeneous machines.

This provides a variety of security benefits for user data, including:

- **Data Distribution** – Data is spread across thousands of systems. As a result, no one system has all of a user's data or all of a company's data. This makes it impossible for an intruder to target and remove a set of systems containing data for any particular customer. It would be like searching for a needle in a haystack.

- **Infrastructure Homogeneity** – Because all systems are the same, security fixes can be very quickly diagnosed and deployed for the entire infrastructure. Google does not need to worry that a particular machine has a different version of the infrastructure software than other systems. Additionally, even if an intruder were to physically breach a datacenter, they would not be able to identify one system from another since they all physically look the same.

- **Failover and Scalability** – Because all systems are the same, any of these systems can be spun up to serve customer data. As a result, the infrastructure can scale and fail over based on dynamic needs.

- **Data Obfuscation** – All user data is stored in a homogeneous Google-proprietary filesystem that does not follow traditional file system storage and access methods (such as NFS or CIFS). As a result, the data is obfuscated and not easily readable by anyone even if they were to breach the system.

Infrastructure for Video ID

Upload Servers

Customers will utilize an SFTP dropbox on specific servers attached to the internet. The login requires a static IP, a public key (sent to YouTube) and a private key (staying at the customer site), and a user acccout. This login is restricted to SFTP only and uses well tested security methods (SSH2, RSA, or DSA). Once logged onto the server, all customers will be separated with a chroot into their sub-directories (and only their subdirectories). The customer can upload multiple video and XML data files into that subdirectory or a child. After uploads have completed, all files are moved into a processing directory and are no longer accessible to the customer. Files will remain on the server for a period of up to 21 days and then are purged.

These servers are separate machines from the streaming servers at YouTube and cannot stream the uploaded files. The machines can only be accessed internally by a limited number of admin account owners.

Database Servers

The database servers receive the files from the upload server (via a private Google network). The videos are transcoded and ID files are created to be used by the Video ID service. Videos are stored indefinitely (320x240 resolution) in the event a new ID file is required in a Video ID upgrade. The videos and ID files are stored under GFS on Google

machines that not accessible via the Internet. Like the upload servers, these servers can only be accessed internally by a limited number of admin account owners.

Google's Own Data on Same Infrastructure

One of the strongest endorsements of Google's security infrastructure is that Google stores our own data on the same infrastructure as our customers. Any security hole would expose critical Google intellectual property and business information, so extreme care and examination was taken to ensure safety and security of the infrastructure.

Internal Security and Change Management Processes

Security is a process that must be a part of the overall culture and operation of the organization. Google takes many measures to ensure that security is central to the process.

Internal Data Access Processes and Policies

Access Policy

LDAP, Kerberos and a Google proprietary system utilizing RSA keys provides Google with secure and flexible access mechanisms. These account mechanisms grant only approved access rights to site hosts, logs, customer information and configuration information. We require the use of unique user IDs, strong passwords, and carefully monitored access lists to ensure appropriate usage of accounts. The granting or modification of access rights ~~are~~ is based on a user's job responsibilities on a need to know basis and must be approved by data owners. Approvals are managed by workflow tools that maintain audit records of all changes.

Furthermore, it is Google's policy to provided system access to individuals who have been trained and require this level of access to perform authorized tasks. Access to systems is logged to create an audit trail for accountability.

Where passwords or are employed for authentication at Google (e.g., login to workstations), password policies that follow best-practices are implemented. These standards include password expiry, restrictions on password reuse and sufficient password strength. For access to extremely sensitive information (e.g., Credit Card data), Google uses hardware tokens.

Code Development Review Process

Design

Major parts of the system and application architecture are documented in a design document before any development has begun. The lead developer will detail the architecture, impact, and security considerations, and circulate amongst the engineering team for open review and approval. Security-focused engineers are involved in the product development process during all phases of the development cycle.

Development and Test

Code change requests as well as system and hardware maintenance are standardized, categorized, and prioritized according to need. To the extent possible, the process and corresponding procedures are documented and designed to drive a controlled framework as well as the proper segregation of duties for the initiation, design, test, approval and migration of changes. The process outlines the change classification and corresponding activities to be performed during each of the phases, which are dependent on the impact the change will have to the system.

The change management process starts with a developer checking out a source code file to make a change. Once development is completed, the developer performs unit tests, if applicable, and a review is performed before the code is checked back into the repository. Google requires that a review independent of the developer be assigned.

Once a file has been properly approved, the release process begins. Code is compiled into a binary, and the binary is transferred to the QA environment where integration testing is performed. Depending on the type of change, dedicated QA resources may exist. If QA resources are unavailable, the lead engineers will take responsibility for performing load and regression testing within the QA environment. Once QA is complete, the binary is moved for migration to production.

Launch

A change is scheduled to be "pushed" to the production environment by the automated change management tool. The push process determines which production files will be migrated by checking the production configuration files which are also managed through the change management process.

Software developers are required to go through a security review when launching any new service on Google infrastructure. During this review, a security engineer from the Google security team will look at the following:

- Review the design document, and review the notes from any previous design review
- Build and run the application or use a test instance of the application to familiarize themselves with the application functionality
- Test against the running application for common known security vulnerabilities
- Review the code for security-sensitive areas such as input validation, file and network I/O, database access, cross-site scripting, and others

The security review is part of the launch checklist process which every application must pass before going into production.

Incident Reporting and Reaction Process

Google employs multiple proactive efforts to monitor for security incidents, including but not limited to inbound security reports, open source alerts, automated perimeter systems, and community alerts. When an Information Security incident occurs, Google security

responds immediately based on the level of threat. Notification of an incident may be generated automatically by monitoring systems or manually by a Google employee. Google works very closely with the security community to track reported issues in Google services and open source tools. More information can be found at http://www.google.com/intl/en/corporate/security.html

When notified of a problem, a Google security engineer makes a risk assessment and begins following prescribed response plans for the issue. Google has documented escalation procedures and communication protocols to address when and how incidents should be escalated as well as who should be notified.

Google continually monitors the production system in a variety of ways such as automated systems that look for predefined events (e.g., router crashes) and the use of statistical dashboards to diagnose and analyze issues (e.g., bandwidth utilization). Thresholds are configured on these monitoring systems so that the system health of network components, servers and other devices can be monitored closely. System reliability teams and customer support technicians respond to alerts generated when the monitoring system detects thresholds have been reached.

Personnel Hiring, Background Check, and Security Training Process

Google has formalized global hiring practices designed to ensure new, rehired, or transferred employees are qualified for their functional responsibility. At a minimum, these practices include verification of the individual's education and previous employment as well as a referral check. Where local labor law or statutory regulations permit, Google may conduct criminal, credit, and/or security checks on all potential employees. The specifics or extent of background checks performed is dependent on the position for which the individual is applying.

Training of personnel is accomplished through the employee's development plan as well as supervised on-the-job training. The development plan is intended to help employees determine which learning activities should be completed to obtain or retain the skills and competencies for their job. This includes any special training necessary for the individual's technical position.

Upon acceptance of employment, all employees are required to execute a confidentiality agreement as well as acknowledge receipt and compliance with Google's Employee Handbook. The confidentiality and privacy of customer information and data is emphasized in the handbook as well as during new employee orientation.

All employees are required to attend security training as part of new hire orientation. At this training, they are instructed about the security policy of the company and escalation procedures.

Every employee has a written job description, and every job description includes the responsibility to communicate timely significant issues and exceptions to an appropriate higher level of authority within the Company.

Data Replication and Data Disposal

Data Replication

Data redundancy is built into the GFS file system, and all data that is written in GFS is replicated at least three times to separate systems. Such protections make sure that a customer's data is protected in the event of a disaster.

Distributed Data Center Architecture

Google does not rely on just one datacenter to run our applications. We operate a geographically distributed set of datacenters to keep services running in the event of incidents and disasters at a single datacenter. Google runs datacenters in over a dozen locations worldwide, and has plans to build several more Google-owned datacenters in the near future. These datacenters are connected via high-speed private links to ensure secure and fast data transfer between datacenters.

Datacenter locations are undisclosed to the public, and data centers are unmarked to ensure optimal data security.

Google's data center management staff is also distributed in multiple geographies to ensure around the clock coverage and system administration that is not location dependent.

Video and ID File Data

"Reference Only" videos are used exclusively for Video ID; the video and ID files for Video ID are in database servers, separate from YouTube video servers. These videos and ID files can be disabled via XML actions. (Note that it is possible to disable the video and still keep the existing ID file active). When either the video or the video-and-ID files are disabled, they become immediately inactivated from the Video ID services. Within 48 hours, disabled ID files are purged and a new Video ID database is fully written across datacenters; this removes all remnants of the ID files. Video files are not deleted from the Video ID servers or backup files.

Data Destruction

Production disks go through a series of data destruction processes when they are removed from our systems. Disks are first logically wiped before they are physically accessed by our production staff. They are then removed from the system and confirmed to be wiped.

Google Privacy Policy

Compliance with Safe Harbor

Google adheres to the US safe harbor privacy principles of Notice, Choice, Onward Transfer, Security, Data Integrity, Access and Enforcement, and is registered with the U.S. Department of Commerce's safe harbor program. This is detailed in the Google privacy policy.

http://www.google.com/privacypolicy.html

version 1.0-01/2008